

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Assistant Commissioner for Patents
Washington, D.C. 20231

Atty. Dkt.: 922-110

Sir:

Date: September 14, 2000

JC913 U.S. PTO
09/662157
09/14/00

Attached for filing is the patent application of:

Inventor: CREEDON et al.

Entitled: **SYSTEM FOR THE DETECTION OF
ASYNCHRONOUS PACKET RATES AND
MAINTENANCE OF MAXIMUM THEORETICAL
PACKET RATE**

and including attachments as noted below:

Declaration, Abstract
 13 pages of specification and claims (including 14 numbered claims), and
 3 sheets of accompanying drawings.
 Record & return the attached assignment to the undersigned.
 Priority is hereby claimed under 35 USC 119 based on the following foreign applications, the entire content of which is hereby incorporated by reference in this application:

Application Number	Country	Day/Month/Year Filed
0012635.9	GREAT BRITAIN	25 May 2000

, respectively.

Certified copy of foreign application is attached.
 Please amend the specification by inserting before the first line --This is a _____ of PCT application _____, filed _____, the entire content of which is hereby incorporated by reference in this application.--
 Priority is hereby claimed under 35 USC 120/365 based on the following prior PCT applications designating the U.S., the entire content of which is hereby incorporated by reference in this application:

Application Number	Country	Day/Month/Year Filed

This application is based on the following prior provisional application(s):

Application No.	Filing Date

respectively, the entire content of which is hereby incorporated by reference in this application, and priority is hereby claimed therefrom.

Please amend the specification by inserting before the first line: -- This application claims the benefit of U.S. Provisional Application No. _____, filed _____, the entire content of which is hereby incorporated by reference in this application.

Verified Statement attached establishing "small entity" status (Rules 9 & 27)
 The Examiner's attention is directed to the prior art cited in the parent application by applicant and/or Examiner for the reasons stated therein.
 Preliminary amendment to claims (attached hereto), to be entered before calculation of the fee below.
 Also attached:

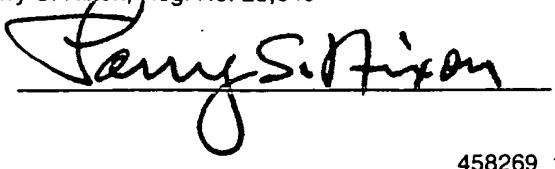
FILING FEE IS BASED ON CLAIMS AS FILED LESS ANY HEREWITH CANCELED

Basic Filing Fee		\$	690.00
Total effective claims 14 - 20 (at least 20) = 0	x \$ 18.00	\$	0.00
Independent claims 4 - 3 (at least 3) = 1	x \$ 78.00	\$	78.00
If any proper multiple dependent claims now added for first time, add \$260.00 (ignore improper)		\$	0.00
	SUBTOTAL	\$	768.00
If "small entity," then enter half (1/2) of subtotal and subtract		\$	0.00
Assignment Recording Fee (\$40.00)		\$	40.00
	SECOND SUBTOTAL	\$	768.00
	TOTAL FEE ENCLOSED	\$	808.00

Any future submission requiring an extension of time is hereby stated to include a petition for such time extension.
 The Commissioner is hereby authorized to charge any deficiency in the fee(s) filed, or asserted to be filed, or which should have been filed herewith (or with any paper hereafter filed in this application by this firm) to our Account No. 14-1140. A duplicate copy of this sheet is attached.

1100 North Glebe Road, 8th Floor
 Arlington, Virginia 22201-4714
 Telephone: (703) 816-4000
 Facsimile: (703) 816-4100
 LSN:ms

NIXON & VANDERHYE P.C.
 By Atty: Larry S. Nixon, Reg. No. 25,640

Signature: 

THIS PAGE BLANK (USPTO)



The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

JCG13 U.S. PRO
09/662157
09/14/00



I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 1 August 2000

**CERTIFIED COPY OF
PRIORITY DOCUMENT**

THIS PAGE BLANK (USPTO)

24 May 2000
Page 1

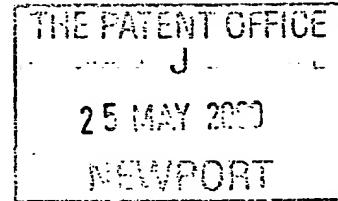
0012635.9

00

25 MAY 2000

Patents Form No. 1/77

The Comptroller
The Patent Office
Cardiff Road
Newport
Gwent NP10 8QQ



REQUEST FOR THE GRANT OF A PATENT

The grant of a patent is requested by the undersigned on the basis of the present application:

1. Title of the invention: SYSTEM FOR DETECTION OF ASYNCHRONOUS
PACKET RATES AND MAINTENANCE OF MAXIMUM
THEORETICAL PACKET RATE

2. Applicants details:

First or only applicant: 3Com Corporation

Country: USA State: Delaware

Address: 5400 Bayfront Plaza
Santa Clara
California 95052-8145
United States of America

639 6741 002

3. Name of agent: BOWLES HORTON

ADP Code: 8805003

Agent's address: Felden House, Dower Mews
High Street, Berkhamsted
Hertfordshire HP4 2BL

4. Agent's reference: 105106

5. The application claims an earlier date under any of Sections 8(3), 12(6), 15(4) or 37(4): NO

6. Declaration of Priority (if any)

Country

Application Number

Priority date

None

24 May 2000
Page 2

Patents Form 1/77 (continued)

7. Inventorship

The applicant(s) is/are the sole inventor or joint inventors: NO

8. Check List of Documents:

A. The application is accompanied by the following number of sheets:

Request: 1

Description: 9

Claims: 1

Abstract: 1

Drawings: 3 (x 2)

B. The application as filed is accompanied by the following:

Patents Form 7/77: NO

Patents Form 9/77: YES

Patents Form 10/77: YES

Priority Document: NO

Translation of priority document: NO

9 Signature



for BOWLES HORTON

SYSTEM FOR DETECTION OF ASYNCHRONOUS PACKET RATES AND MAINTENANCE OF MAXIMUM THEORETICAL PACKET RATE

Field of the Invention

5 This invention relates to systems, and particularly multiple chip devices, for the handling of data in the form of data packets which are separated by inter-packet gaps and are caused to proceed at a local clock rate and particularly to systems wherein data packets are transferred from one clock regime to another, as for example between one ASIC (application specific integrated circuit) and another. The invention also relates to the maintenance of a maximum theoretical rate of transmission of packets by the addition and removal of non-packet words (usually idle or preamble bytes).

10

Background to the Invention

15 Packet handling devices, such as switches (both bridges and routers) repeater, hubs and the like are well known in the art of communication systems which employ data packets, such as Ethernet packets, for the conveyance of data from user to user. It is now customary to realise complex devices in the form of a multiplicity of interconnected 'chips' i.e. application specific integrated circuits which typically embody the physical layer devices (PHYs), media access control (MAC) devices, look-up engines, control registers and some, though not usually all, memory for the temporary storage of packets in the interval between the reception of packets by the device and their dispatch from a port or ports of the device. The invention is generally applicable in circumstances where a given device is constituted by more than one application specific integrated circuit.

20

25

In most packet-based communication systems, the data stream comprises packet data and words (usually bytes) which are of practical necessity but do not comprise message information. For example, in current Ethernet systems, there is a prescribed inter-packet gap

constituted by 12 bytes, followed by 7 bytes of a 'preamble' and a 'start of frame delimiter' (SFD) byte (a total of 20 bytes), followed by the packet itself. The maximum theoretical packet rate is achieved when transmitting minimum-sized packets with a minimum allowed inter-packet spacing. The present invention concerns the maintenance of such a rate notwithstanding slight discrepancies in controlling clocks, for example between a media access control (MAC) device in one chip and a physical layer device (PHY) in a different chip which has its own crystal-controlled clock source.

It is obvious that a single device constituted by a plurality of ASICs should have in common a high precision clock frequency. In practice ASICs of this general nature require a set or suite of clock frequencies derived from or related to a master or top rate clock frequency. As a practical matter, different ASICs commonly have their own respective clock signal generating circuits and although theoretically the clock frequencies should be exactly the same, in practice there are small variations. This arises in particular when two such clock signals are derived from different crystal sources designed to have the same nominal frequency. In a system where there are or may be significant differences between clock frequencies on respective ASICs, it is necessary to employ at, for example, the receiving end of a link between ASICs an elasticity buffer into which data packets received over the link are written, usually one byte at a time, and read out at the local clock rate.

1 (1)

15

20

The present invention more particularly concerns operation wherein two adjacent clock domains have only slightly different clock frequencies, the difference being for example less than a few hundred parts per million.

۲۵

If the clock in the receiving chip (e.g. the PHY) is slightly slower than that in the source (e.g. the MAC) there is significant danger that the 'elasticity buffer' will be overrun. Various expedients to avert this danger, on the assumption that the proximity of the danger can be detected. Those which involve discarding whole packets are unsuitable for high performance systems and those which rely on high and low watermarks of a FIFO tend to require large

FIFO's and unnecessary complexity. It is also possible to discard preamble bytes but by itself this is an inappropriate solution, because some devices in the network may require the full complement of preamble bytes for correct operation.

5 Summary of the Invention

A feature of the invention concerns the detection of small differences between clock frequencies employing an elasticity buffer have a write process controlled by one clock and a read process controlled by a second clock. According to one aspect of the invention a buffer 10 having a minimal number, and more particularly five recycling storage locations is employed to detect small differences in clock rates and to generate a resynchronisation command, which may indicate a need to discard a byte (from a preamble or inter-packet gap).

In a further aspect of the invention the resynchronisation command is employed to cause at an 15 earlier stage (e.g. a MAC chip) the insertion of an extra preamble or idle byte pertinent to an inter-packet gap. The detection of the inserted byte may initiate the resynchronisation of the elasticity buffer and the subsequent discarding of the inserted byte.

Further features of the invention will be apparent from the following detailed description with 20 references to the accompanying drawings.

Brief Description of the Drawings

Figure 1 is a schematic representation of a minimal elasticity buffer employed to detect minor 25 variations in clock frequencies.

Figure 2 is schematic diagram of an interface between application specific integrated circuits, employing an elasticity buffer as described with reference to Figure 1.

Figure 3 is a diagram illustrating the operation of the elasticity buffer.

Detailed Description

5 Figure 1 of the drawings illustrates part of an interface between two domains, which have nominally identical clock frequencies which in practice may vary slightly, the range of variation being in the present invention not more than a few hundred parts per million.

10 The interface 10 shown in Figure 1 has an input link, typically a multiplicity of lines on which for a given clock cycle the respective bits of a byte are transmitted, for example as described in published British patent applications numbers 2336074 and 2336075. As described in those earlier applications the eight data lines may be accompanied by a ninth or control line which may be employed for phase alignment of the signals on the various data lines.

15 It is presumed that data from a first ASIC (e.g. a MAC) is transmitted by a set of lines 11 in parallel form and successive bytes are stored in successive locations a, b, c, d and e of an elasticity buffer 12, being read out also in a cyclic fashion on data lines 13. Writing of input packets into the elasticity buffer is controlled by a write pointer 14, i.e. a set of signals which determine which location is selected for writing the byte, under the control of a write pointer generator 15 itself clocked by a clock in domain A. Bytes are read out from the elasticity buffer 12 in a cyclic sequence by means of a pointer 16 produced by a read pointer generator 17 controlled by a clock in domain B. The domain for clock A is termed herein the write clock domain whereas the domain for clock B is termed herein read clock domain. The write clock domain will be that pertaining to the ASIC from which the data has come and the read clock domain will be that of the ASIC which includes, or receives data from, the elasticity buffer 12.

The read clock domain includes a slip-detect circuit 18 which responds to the pointers 14 and 16 to determine the state of relative synchronism of the write clock and the read clock.

A five-deep elasticity buffer represents the smallest feasible buffer in the presence of jitter and small variation in the clock rate. Correct synchronised operation of the buffer requires the write pointer to be, for example, loading location *a* while the read pointer unloads location *d*.
5 If read pointer is unloading location *c* or *e* in the same cycle as location *a* is loaded then jitter is present and the read clock is early and late relative to the write clock respectively. If it unloads location *a*, in the same cycle as location *a* is loaded, then synchronisation is required, the read clock being too early. If read clock unloads location *b* in the same cycle as location *a* is loaded, the read clock is too late and resynchronisation is needed.

10

It would be possible to employ an elasticity buffer with even fewer locations, specifically three, if corruption of the data stream were tolerable.

15 Figure 3 illustrates (i) the write pointer (which toggles in sync with its controlling clock) and the positions of the write pointer for a succession of cycles of the write pointer clock (ii) the variation of an extra bit (wr3) of the write pointer (iii) the variation of three bits of the write pointer, (iv) the read pointer (which toggles in sync with its controlling clock) and (v) the significance of the various positions of the read pointer for the various positions of the write pointer. The extra bit is added to the 'count of five' write counter, which effectively counts to
20 ten.

25 The block 18 compares, in the read clock domain, three bits of the read pointer with a single (synchronised) bit of the write pointer. These are sufficient to generate a signal indicating the occurrence of the positions that anticipate a need for the buffer to resynchronise, the read pointer relative to the write pointer and to cause discard of the byte. The toggle of the upper (wr3) bit will indicate an exact position of the write pointer once every five cycles, which is sufficiently frequent to provide an opportunity to resynchronise during a preamble (eight clock cycles) or an IPG gap (twelve clock cycles). The use of this single write bit means that only one bit has to cross the boundary between clock domains. It would be possible to

perform the comparison in the write clock domain, employing the most significant bit of the read pointer and three bits of the write pointer, performing resynchronisation by controlling the write pointer, leaving the read pointer free running.

5 Figure 2 represents a system in which the subsystem 10 shown in Figure 1 is employed in the maintenance of a maximum packet rate. A first ASIC 21 (e.g. a MAC) includes a system memory 23 and a memory control 24. Packets are read out from system memory at a rate determined by the clock frequency in clock domain A and conveyed by means of link 25 to ASIC 22 (e.g. a PHY). Such packets are received by the elasticity buffer 12 (Figure 1), being written into and read out of the buffer as described with reference to Figure 1.

10

15 Figure 2 includes a packet monitor and control block 25 which is coupled across the clock domains to MAC 24 by means of a line 26. This line carries a one bit 'AddOneIdle' signal that toggles (i.e. changes from 0 to 1 or 1 to 0) when the MAC is instructed to add an idle byte to the data stream on line 11. The block 25 receives a signal 'ResyncAddDone' from subsystem 10, indicating that the buffer 12 has performed a resynchronisation and added an idle (by reading the same location twice). This signal is intended only for monitoring purposes. Block 25 receives a signal 'ResyncDropRqst' indicating that the buffer should discard a byte (by moving on two locations instead of one). In response to this signal block 20 25 toggles the signal on line 26. Block 25 also receives a signal (for monitoring purposes) denoted 'ResyncDropDone' when the buffer has discarded a byte.

25 The PHY also monitors for an indication (explained below) of the extra byte which has been added to the data stream by MAC 24 and in response to that indication causes the resynchronisation of the read pointer to drop the relevant byte.

The operation of the system is based on (a) the monitoring of the phase relationship of the read clock and write clock in terms of the positions of the read and write pointer; (b) the signalling by the slip detector (18) of the need to resynchronise; (c) the signalling of the

upstream device to insert an extra byte; and (d) the detection of the extra byte and the consequential discard of the byte by the resynchronisation of the read pointer. In this manner the discrepancy between the clocks is compensated by an operation which will cause discard of a byte but an additional byte is inserted in anticipation of the discard.

5 One advantage of the invention is that the anticipation of the need to drop a byte allows the transmitter, in this example the MAC chip, sufficient time to prepare for and effect the insertion of the compensatory byte, at a position which is at the choice of the designer. The maximum time that the transmitter can delay in inserting the byte is a function of the difference in clock frequencies, the packet sizes and the size of the elastic buffer. In practice there are likely to be two or three packet times available.

10 In more detail, in the present example, the receiver, that is to say the PHY by monitoring the progression and the positional phase relationship of the read and write pointers will detect the need to resynchronise, particularly by advancing the read pointer by two locations, in a way that will have the effect of dropping a byte. In consequence, as indicated above, the 'AddOneIdle' signal is toggled. Although it would be possible to adopt some other form of indication, for example by sending a signal which is active high and deasserting that signal on reception of an acknowledgement from the transmitting MAC, the advantage of toggling the 15 'AddOneIdle' signal is that it is simple and it is not dependent on any circuitry which has to produce, for example, a signal which is long enough to be detected in a different clock domain.

20 The transmitter, the MAC 24, responds to the toggling of the 'AddOneIdle' signal and, in a manner which is simple in itself, increases the number of bytes between data portions of packets. In an Ethernet system wherein the inter-packet gap is 12 bytes, the inter-packet gap may be increased to 13 bytes. Alternatively, the preamble may be increased from 7 to 8 bytes. A MAC device may for example perform this insertion as part of the read out process of 25 packet data from the system memory 23.

5 It is desirable for the MAC to indicate the existence and preferably the location of the added byte. This may be performed by any one of a variety of techniques. It could simply add a signal which is coded as a 'extra byte' signal, so that it is recognisable as a superfluous byte just as an 'idle' byte is so recognised. Alternatively, it could encode such information with other packet control signals so that a 'extra byte tag' is aligned with the specific byte involved. The interfaces which are described in the aforementioned patent applications provide a control line in parallel with a multiplicity of data lines and the extra byte identifier may be carried on such a control line if desired. As a further alternative, the packet itself may 10 be tagged in its header to indicate that it is preceded by a preamble of more than the usual length.

15 Finally, the PHY notes the 'extra byte' or the control code or tag indicating it, and as it extracts bytes from the elasticity buffer moves (by means of resync circuit 20) the read pointer one position away from the direction of the write pointer, thereby dropping or discarding the required byte. The extra byte may be detected by appropriate control logic in the elasticity buffer, which logic is in a state awaiting the extra byte, because it has asserted 'ResyncDropRqst'. The byte is dropped by the resynchronising of the elasticity buffer and the signal 'ResyncDropDone' is toggled to indicate that the byte has been dropped.

20 The net effect is that the packet rate is identical to that produced by an ideal transmitter operating at the clock frequency of the PHY chip, at the maximum theoretical rate within the tolerance allowed.

25 When the PHY crystal is slightly faster than the MAC crystal, that is to say the signal in clock domain B is slightly faster than the signal in clock domain A, there is no need for the PHY to communicate with the MAC. The elastic buffer itself may detect when it is close to the point where the read pointer catches up on the write pointer, that is to say within one location of it, and that may be used during, for example, the inter-packet gap, to cause the reading of the

same location twice, which will effectively add an idle byte while resynchronising the read pointer.

5

10

15

20

25

Claims

1. A system for transferring a data stream comprising data packets separated by non-packet words from a first clock domain to a second clock domain, the clock domains having similar but not necessarily identical clock frequencies, comprising an elasticity buffer, means for writing the data stream into the elasticity buffer in a cyclic sequence under the control of the clock frequency in the first clock domain, reading the data stream out of the elasticity buffer in a cyclic sequence under the control of the clock frequency in the second domain, means for monitoring the two sequences to provide an anticipatory signal indicating that the reading sequence approaches proximity to the writing sequence, means in the first clock domain for inserting a non-packet word into said data stream, and means in the second clock domain for detecting the existence of the inserted non-packet word and for causing the buffer to advance the reading cycle thereby to discard the said inserted non-packet word.
2. A system according to claim 1 wherein the inserted non-packet word is an idle or preamble byte.
3. A system according to claim 1 or claim 2 wherein the transmitter means is a media access control device and said receiver means is a physical layer device, these two devices having respective clock signal sources of the same nominal clock frequency.
4. A system according to any foregoing claim wherein the elasticity buffer has not more than five storage locations.
5. A system according to claim 4 wherein the means for writing and the means for reading include a write pointer and a read pointer respectively.
6. A system according to claim 6 wherein the means for monitoring includes a slip detector responsive to the write and read pointers.

Abstract

5 A system for transferring a data stream comprising data packets separated by non-packet words from a first clock domain to a second clock domain includes an elasticity buffer into which the data stream is written in a cyclic sequence under the control of the clock frequency in the first clock domain and from which the data stream is read out in a cyclic sequence under the control of the clock frequency in the second domain. The two sequences are
10 monitored to provide an anticipatory signal indicating that the reading sequence approaches proximity to the writing sequence. A non-packet word is inserted into the data stream in the first domain. In the second clock domain the existence of the inserted non-packet word is detected and the buffer is caused to advance the reading cycle thereby to discard the said inserted non-packet word.

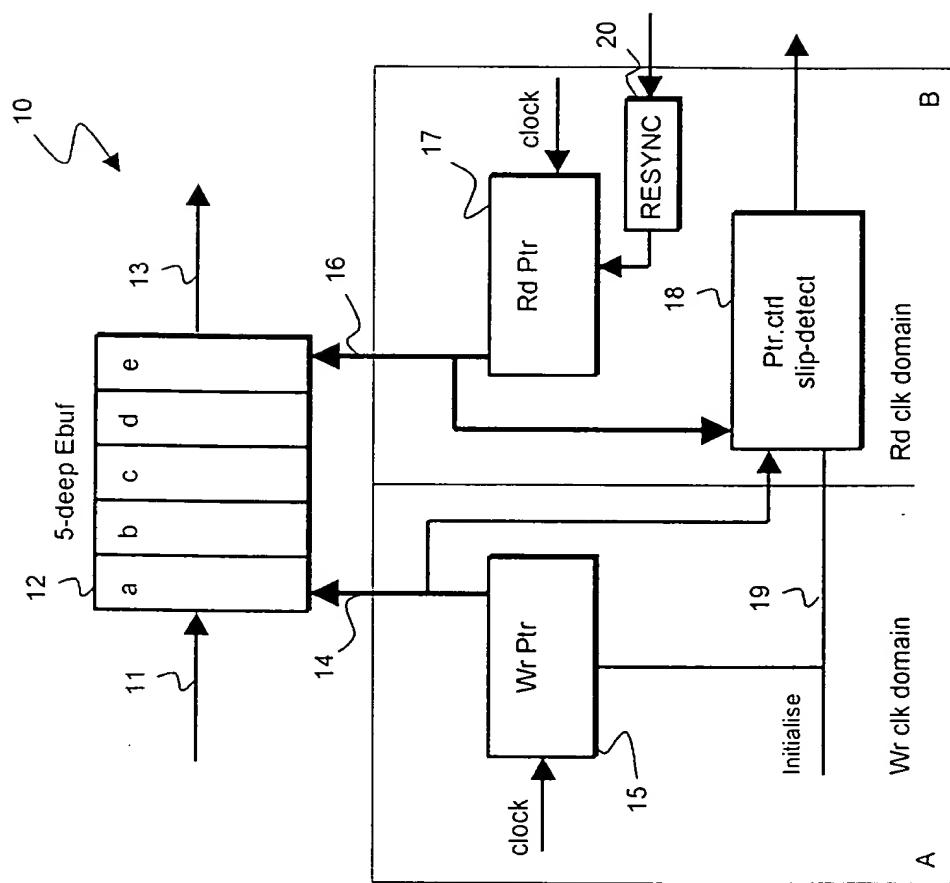
15

20

25

THIS PAGE BLANK (USPTO)

FIG.1



THIS PAGE BLANK (USPTO)

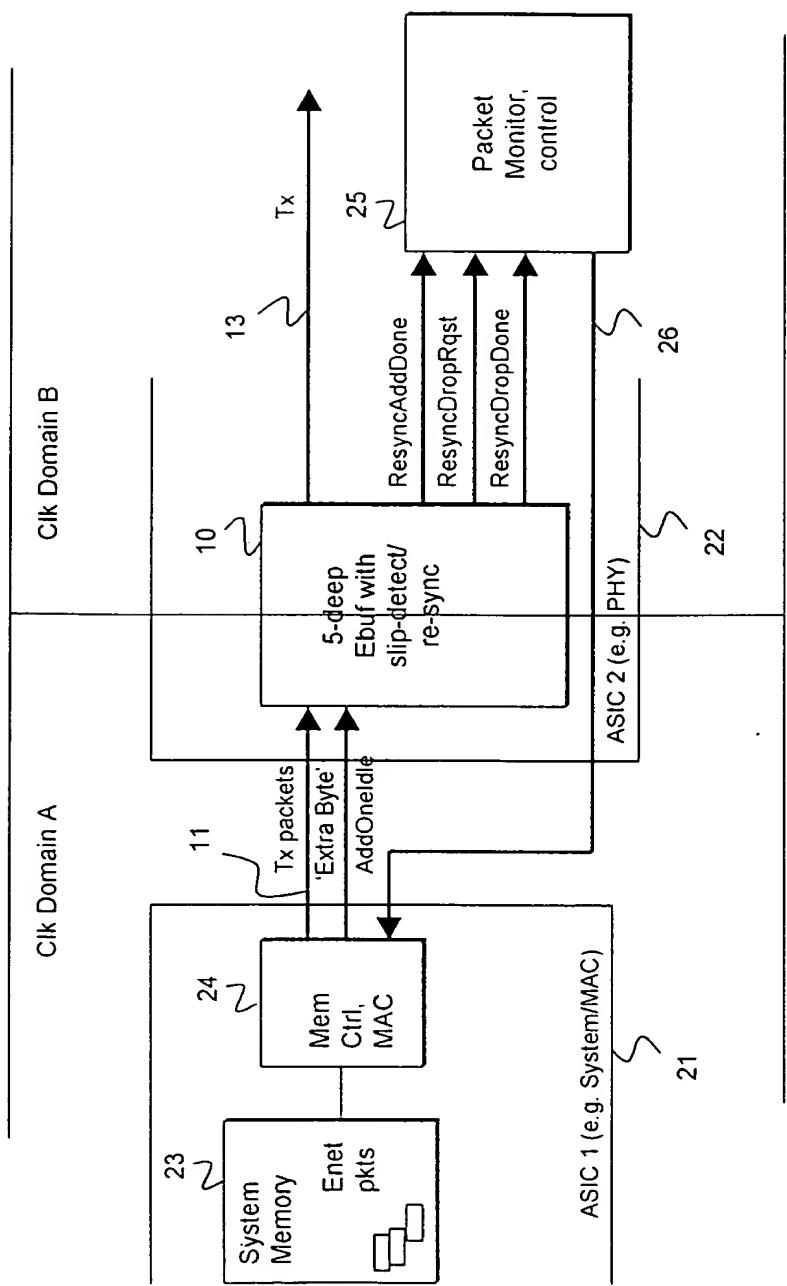
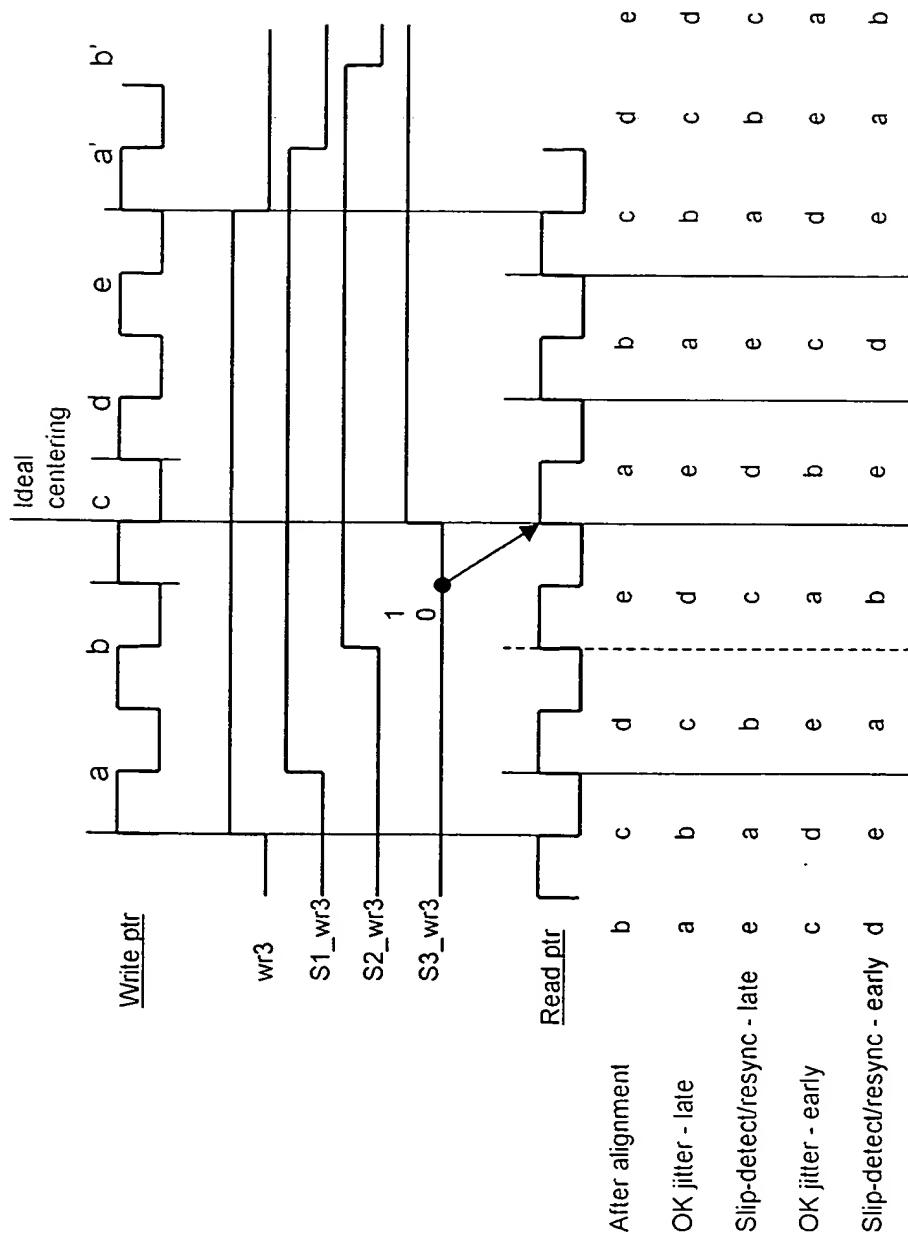


FIG.2

THIS PAGE BLANK (USPTO)

FIG.3



THIS PAGE BLANK (USPTO)